

**REAL-TIME MONITORING OF PERCEIVED QUALITY OF PACKET VOICE
TRANSMISSION**

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This is a continuation-in-part of U.S. Patent Application Serial No. 09/778,186 filed on February 7, 2001, which is a continuation-in-part of U.S. Patent No. 6,370,120 the content of which is relied upon and incorporated herein by reference in its entirety, and the benefit of priority under 35 U.S.C. §120 is hereby claimed. The present application also relates to provisional U.S. Patent Application No. 60/492,597, filed August 5, 2003, which is incorporated by reference herein in its entirety and to which the benefit of priority under 35 U.S.C. §119(e) is hereby claimed.

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] The present invention relates generally to telecommunications networks, and particularly to monitoring the quality of packet voice transmissions in real-time.

2. Technical Background

[0003] New telecommunications technologies are emerging that employ packet switching instead of the traditional circuit switched technologies provided by the public switched telephone networks (PSTN). These packet switched technologies are being used to support the transmission of digitized voice signals over a data network such as the Internet. Providing telephone-like full duplex voice over an Internet Protocol network (VoIP) is particularly important. VoIP services are attractive to commercial long-distance carriers because they enable the use of global Internet transport facilities to carry traffic that is presently being carried over dedicated circuit switched facilities. The potential benefits of VoIP are enormous in terms of better utilization of network bandwidth to support telephone traffic and the economies of scale from the use of one kind of transport for all telecommunications services. However, a major impediment to the immediate adoption of VoIP services relates to the user perception of the quality of voice communications using VoIP.

[0004] Traditional telephone connections have been subject to impairments in the form of noise, attenuation, distortion, cross-talk and echo. Such impairments are particularly common in analog portions of the network, such as subscriber loops and frequency domain multiplexing equipment. Digital transmission alleviates many of these problems but also introduces quantization noise and distortions due to bit errors in the digital signal. Even with perfect digital transmission applied to long haul transmissions, a typical telephone connection still includes many analog components where impairments can occur.

[0005] A poor connection or malfunctioning piece of equipment can produce conditions that a telephone customer will find objectionable or intolerable. When there is a high incidence of poor connections, customers may complain to the service provider or to a regulatory authority, or simply change long distance carriers. Thus, the perceived quality of a service provider's service is a major factor affecting the reputation and marketability of long distance telephone services.

[0006] To guard against poor quality, service providers have developed methods to obtain objective quality measurements upon a line, piece of equipment, or an end-to-end telephone connection. These measurements can help the service provider detect and gauge impairments, pinpoint weak elements, and correct deficiencies that degrade user perception of quality. The effects of extreme fault conditions on user perception of quality is clear. There are easily discernable thresholds for "no effect" and "substantial degradation" conditions. As a result, the average consumer has to come to expect a certain quality of service from the PSTN.

[0007] With the proliferation of voice-over-packet technologies, maintaining a quality of service comparable to the PSTN is a major concern of service providers, equipment vendors, and ultimately the consumers of packetized telecommunications services. Unlike circuit switched traffic, real time voice transmission using packet switched technologies is sensitive to packet loss, packet delay, and packet jitter occurrences which are characteristic of packet switched networks. Packet loss and packet delay variations may impact the ability of a voice codec to faithfully reproduce a digitally encoded voice signal. When a received packetized voice transmission is

missing some packets, a codec may provide an audio signal that is distorted, garbled, or otherwise degraded.

[0008] In one approach that has been considered, the IntServ and Diffserv protocols have been proposed for improving the reliability and consistency of packet transport. (For reference, the IntServ and Diffserv approaches are described in documents, RFC 1638 and 3317 respectively, promulgated by the Internet Engineering Task Force (IETF).)

[0009] The Real Time Control Protocol (RTCP) has also been considered for obtaining real-time measurements of the receipt of packet data, and for reporting the measurements to a sender or to a network quality monitoring location. (For reference, RTCP is described in IETF document RFC 1889 and in ITU Recommendation H.225.0.)

[0010] One drawback to these various approaches is that, while the reporting of packet arrival statistics provides some estimate of data transmission quality, there is no consideration of the extent to which packet loss, packet delay, and packet jitter affect the perceived quality of a reconstructed voice signal.

[0011] Another drawback relates to the fact that the manner in which data transmission quality affects perceived quality of a voice channel is often dependent on the coding scheme employed by the codec. Various codec schemes may exhibit differing susceptibilities to packet-transmission variations and the dependency may be quite non-linear for some codec schemes. For example, a given packet loss rate or jitter may have very little effect on a G.711 encoded speech signal, whereas an identical packet loss rate or jitter may seriously degrade a more sophisticated G.723 encoded signal. Even if the particular codec scheme is known, an accurate model of codec behavior is required to map the observed packet characteristics to a perceived quality level. The introduction of a new standard encoding scheme, or a proprietary encoding scheme or encryption scheme, would necessitate the development and deployment of new models.

[0012] For these reasons, the mere reporting of packet loss statistics is inadequate. Without accurate information about the perceived quality level being experienced by

users, a service provider is not certain when corrective action is necessary or what corrective action needs to be taken to improve packet transmission performance. A service provider may shift traffic or resources unnecessarily to improve quality of service (QOS) when, in fact, the perceived level of degradation is negligible to the users of the service. Likewise, a seemingly minor packet delivery problem might be causing big problems in user perception due to the particular codec scheme being used.

[0013] What is needed is a means for automatically monitoring and reporting in near real-time the quality of a packet voice transmission as perceived by a user receiving voice communications.

SUMMARY

[0014] The present invention provides a system and method for automatically monitoring and reporting, in real-time, the perceived quality of a packet voice transmission. The present invention is versatile in that it may be configured to report a simple raw distortion measurement or provide sophisticated estimates of perceived voice quality, using MOS and P[UDI]. In the latter case, the present invention provides an analytical mechanism for detecting when packet-switched voice services are likely to be perceived by users to be noticeably inferior to familiar PSTN services. This embodiment of the present invention may be based on the VQES model described above. The present invention also provides a network performance data collection and analysis system for packet switched networks. Further, the system of the present invention functions to monitor voice quality of VoIP services to support timely detection and notification of indications of deterioration.

[0015] One aspect of the present invention includes a method for monitoring perceived quality of a packet-switched voice service in a network. The method includes the step of receiving a packetized voice communication via the packet-switched voice service. At least one objective measurement is obtained from the received packetized voice communication. An estimate of user perceived quality of voice data is derived from the at least one objective measurement. The user perceived quality of the received

digital voice data is provided to a far end user or a service provider. The steps of receiving, obtaining, deriving, and providing are performed in real-time. As taught herein, the objective measurement may be obtained from a digital signal output from a receiving codec such that the impact of transmission impairments upon the codec scheme are inherently taken into account in the measurement.

[0016] In another aspect, the present invention includes a system for monitoring the quality of a packet-switched voice service. The system includes a memory element configured to store a reference matrix that mathematically models likely user perception of acceptable quality of voice service. The reference matrix is empirically derived from a plurality of objective voice measurements. A measurement probe is configured to obtain a plurality of test measurements for each call placed over the packet-switched voice service. A computer is coupled to the memory element and the measurement probe. The computer is programmed to derive a test matrix from the plurality of test measurements. The computer processes the test matrix, in near real time, to determine the quality of voice over the packet-switched voice service. The quality of service is determined by comparing the reference matrix to the test matrix.

[0017] In another aspect, the present invention includes a method for monitoring the quality of a packet-switched voice service in a computer system having a graphical user interface, the graphical user interface including a display and a selection device. The method includes receiving an alarm signal from the computer system. The alarm signal is generated in response to determining that the quality of voice over the packet-switched voice service is below a predetermined level. A message is displayed in response to the alarm signal. An amplifying display icon is selected with the selection device. A human readable description of the alarm signal is displayed in response to the step of selecting. At least one indicator of likely user perception of the quality of voice carried over the packet-switched voice service is displayed.

[0018] In yet another aspect, the present invention includes a system for monitoring the quality of a packet-switched voice service in a network. The system includes a measurement device configured to obtain at least one objective measurement from a packetized voice communication in real-time. A processor is coupled to the

measurement device. The processor is configured to derive user perceived quality of voice data from the at least one objective measurement and provide the user perceived quality of voice data in real-time.

[0019] In yet another aspect, the present invention includes a method for monitoring the quality of a packet-switched voice service. The method includes providing a reference matrix that mathematically models likely user perception of acceptable quality of voice service. The reference matrix is empirically derived from a plurality of objective voice measurements. A plurality of test measurements are obtained for each call placed over the packet-switched voice service. A test matrix is created from the plurality of test measurements. The test matrix is processed, in near real time, to determine the quality of voice over the packet-switched voice service. The quality of service is determined by comparing the reference matrix to the test matrix.

[0020] Additional features and advantages of the invention will be set forth in the detailed description which follows, and in part will be readily apparent to those skilled in the art from that description or recognized by practicing the invention as described herein, including the detailed description which follows, the claims, as well as the appended drawings.

[0021] It is to be understood that both the foregoing general description and the following detailed description are merely exemplary of the invention, and are intended to provide an overview or framework for understanding the nature and character of the invention as it is claimed. The accompanying drawings are included to provide a further understanding of the invention, and are incorporated in and constitute a part of this specification. The drawings illustrate various embodiments of the invention, and together with the description serve to explain the principles and operation of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0022] Figure 1 is a diagram of a hybrid telecommunications network in accordance with the present invention;

Figure 2 is a diagram of an IP telecommunications network in accordance with the present invention;

Figure 3 is a block diagram of a system for evaluating the perceived quality of a packet switched voice channel in accordance with an embodiment of the present invention;

Figure 4 is a chart showing a process for providing a reference matrix in accordance with the present invention;

Figure 5 is a diagram illustrating the semantic content of a test matrix, a reference matrix and a reference pattern matrix in accordance with an embodiment of the present invention; and

Figure 6 is a block diagram of a system for evaluating the perceived quality of a packet switched voice channel in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0023] Reference will now be made in detail to the present exemplary embodiments of the invention, examples of which are illustrated in the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts. An embodiment of the network of the present invention is shown in Figure 1, and is designated generally throughout by reference numeral 10.

[0024] The present invention is directed to a method for monitoring perceived quality of a packet-switched voice service in a network. The method includes the step of receiving a packetized voice communication via the packet-switched voice service. At least one objective measurement is obtained from the received packetized voice communication. User perceived quality of voice data is derived from the at least one objective measurement. The user perceived quality of voice data is provided to a user. The steps of receiving, obtaining, deriving, and providing are performed in real-time. Thus, the present invention provides a system and method for detecting when packet-switched voice services are likely to be perceived by users to be noticeably inferior to familiar PSTN services.

[0025] In an approach developed by one of the present inventors, a Voice Quality Evaluation System (VQES) has been used to evaluate the quality of network services. The VQES is a theoretical model used to estimate indicators of likely user perception of quality of voice carried over a telephone connection, from measurements describing the quality of the voice connection. The inputs to the VQES are measurements of, or nominal values for, seven measures of a telephone connection: received speech signal power, or test signal power attenuation; C-message noise; echo path loss and echo path delay; characteristics of natural speech waveforms; dropped packet rates; and round trip signal transport delays. All of these objective characteristics are measured in a way that serves as a reliable indicator of when users will complain of distortion, low volume, echo, noise, or a noticeable delay in the flow of conversation. The resultant data sets are transformed into measures that can be readily compared to determine whether users are likely to perceive a difference in voice quality among competing services tested. Measures of this kind produced by the VQES include, for example, estimates of MOS and P[UDI]. The MOS is a Mean Opinion Score (MOS) that is derived from subjective user evaluation of voice quality under the conditions reflected in the distribution of the objective measurements described above. P[UDI] refers to the proportion of calls in the sample for which the subjective user evaluation of the effects of the impairments would be that they rendered the call "unusable", "difficult", or "irritating." As embodied herein, and depicted in Figure 1, an example of a hybrid telecommunications network in accordance with the present invention is disclosed. Hybrid telecommunications network 10 includes public switched telephone network (PSTN) 14 coupled to packet switched network 18 by way of gateway 16. PSTN 22 is also coupled to packet switched network 18 by way of gateway 22. In the example shown in Figure 1, standard POTS telephone 12 initiates a call with standard POTS telephone 24 by way of circuit switched PSTN 14. PSTN 14 routes the call to gateway 16. Gateway 16 converts time division multiplexed (TDM) signals into packets that are compatible with packet switched network 18. Gateway 16 has access to a directory of IP addresses of exit gateways. Gateway 16 uses the directory to select gateway 20. Packet switched network 18 transports the packets to gateway 20. Gateway 20 converts the packets back into a TDM format compatible with PSTN 22. PSTN 22 routes the TDM signal to telephone 24.

[0026] Packet switched network 18 is commonly referred to as a converged network. A converged network is adapted to carry various types of media, such as voice, data, and audio or video streams. Each of these media types are digitized and formatted into discrete packets of data. Packet switched network 18 may be of any suitable type, but there is shown by way of example an Internet Protocol (IP) network adapted to carry packet-switched voice, or Voice-over- IP (VoIP). In one embodiment, network 18 is the public Internet. The Internet is the term used to describe the worldwide collection of interconnected networks that are linked together by the use of the Internet Protocol (IP). This worldwide collection of networks includes a host of networks sponsored by academic, commercial, governmental, and military entities. In another embodiment, packet switched network 18 may include a packet switched network, or a collection of packet switched networks interconnected by a common protocol, not connected to the Internet. An example of such a network would be a privately owned intranet, or enterprise network. In yet another embodiment, network 18 may include any or all of enterprise networks, intranets, and the public Internet.

[0027] As embodied herein, and depicted in Figure 2, another example of a telecommunications network in accordance with the present invention is disclosed. In this example, network 10' is an end-to-end packet switched network. IP telephone 12' generates packet switched voice signals that are provided directly to packet switched network 18'. The packets are routed to IP telephone 24' by network 18'. Those of ordinary skill in the art will recognize that IP telephone 12' and IP telephone 24' may include a standard telephone that is coupled to packet interfacing equipment.

[0028] Further details on hybrid and end-to-end packet-switched networks are provided in commonly assigned patent applications Serial No. 08/751,023 filed November 18, 1996 (Attorney Docket No. VON-96-001) and Serial No. 08/798,350 filed February 10, 1997 (Assignee Docket No. VON-97-004), each of which is incorporated herein by reference.

[0029] The present invention provides a real-time, or a near real-time, monitoring system for packet-switched telephone networks. The system works equally well with the hybrid network shown in Figure 1 or the end-to-end packet-switched telephony system

shown in Figure 2. Those of ordinary skill in the art will recognize that both Figure 1 and Figure 2 are simplifications of what a "real world" network would comprise. Much equipment is not illustrated for the sake of simplicity and ease of illustration. For example, the routing within the packet switched networks is performed by numerous routers as well as other components.

[0030] As embodied herein and depicted in Figure 3, a system 300 for evaluating the quality of a packet switched voice channel is disclosed. Measurement probe 310 is coupled to the packet-switched voice channel (not shown). Probe 310 is coupled to measurement circuitry 320. Measurement circuitry provides computer 340 with test measurements which include objective characteristics of the voice signal.

[0031] In one embodiment the objective characteristics are those reflecting or determining the two principal ways that packet-switching can degrade perception of voice quality. These comprise: effects of dropped packets on user perception of speech distortion, and effects of additional round trip delay due to packet latency on natural conversational rhythms. The objective measures used for illustration here are the dropped frame rate and the round-trip packet latency. These are, however, for illustration only, as those with ordinary skills in the art of voice quality measurement will be readily able to identify alternatives that will accomplish the same end without deviating from the present invention.

[0032] In another embodiment, probe 310 and measurement circuitry 320 are implemented in selected routers disposed in the network. In another embodiment some of the computing functionality may also be disposed in the selected router. For example, Cisco Systems, Inc. manufactures a series of router devices commonly referred to as the Gigabit Switch Router 12000 Series of routers. These devices are commonly employed in the backbone of the Internet. The 12000 Series uses 200 MHZ R5000 processors. The largest model in the 12000 Series is the 12016 model, which accommodates up to 320 G/bits of data per second and can process up to 60 million data packets per second. In other embodiments, these elements are implemented in gateways or in other network elements. Alternatively, probe 310 and circuitry 320 are implemented in a stand-alone piece of equipment that is not part of the network.

[0033] The objective characteristics measured by circuitry 320 include dropped frame rate (DFR) and round trip packet-switched delay (RTL). As described earlier, each of these objective characteristics affect one or more of a plurality of measures describing the quality of the voice signal as perceived by users. Using techniques described below, measurements of the objective characteristics can therefore be translated into quantifications for each of the subjectively assessed quality characteristics.

[0034] System 300 also includes a database 330 that stores a reference matrix. The reference matrix provides reference quality information for voice signals. The reference matrix will be described in more detail below. Database 330 is coupled to computer 340. Computer 340 runs software that estimates the likely user perception of quality information for a given input voice signal. Once the quality estimates are calculated, they are utilized by maintenance display 350. The functionality of display 350 is described in more detail below. However, in one embodiment, display 350 is coupled to real-time control circuitry. In another embodiment, display 350 provides service messages to a technician in a network maintenance location. Service messages could also be transmitted via pagers, e-mail, audio display and/or visual display. Those of ordinary skill in the art will recognize that database 330, computer 340, and display 350 may be co-located in a network maintenance facility 30.

[0035] Essentially, the present invention operates as follows. A reference matrix that mathematically describes a totality of conditions for which users are likely to find quality of a voice service acceptable is stored in database 330. As will be described in more detail below in conjunction with Figure 4, the reference matrix is empirically derived from the seven objective voice characteristics mentioned above. After the test measurements are obtained by probe 310 and circuitry 320, computer 340 creates a test matrix using the measurements for a plurality of calls obtained by circuitry 320. The test matrix is processed, in near real time, to determine whether the quality of voice over the packet-switched voice service for the time period of the observations was acceptable. This determination is made by comparing the reference matrix to the test matrix. An alarm is generated if the results of the comparison of the reference and test matrices indicates the possibility of unacceptable quality. The alarm notification is accompanied by a reference pattern matrix providing an indication of the matrix cells for which the

comparison was negative. Information indicating the underlying cause and the proper response is displayed for system operators. A description of each of the above mentioned processing routines is provided below.

[0036] Referring to Figure 4, a process for providing the reference matrix is described. In particular, data flow diagram 400 shows the manner in which data from objective and subjective measurements of a communications service may be processed in accordance with the present invention. Data flow diagram 400 initially employs non-packet switched data 410. Non-packet switched data 410 is derived by transforming objective measurements 412 into perceived impairment data 416 by the application of transform 414. Objective measurements 412 include signal power (PWR), C-weighted noise (NSE), echo path delay (EPD), echo path loss (EPL), and waveform distortion (DST). Impairments 416 are analogous to subjective ratings of perceived impairments, namely low volume (LV), noise (NS), echo (EC), and speech distortion (SD). Note that various combinations of objective measurements 412 will affect different ones of the subjectively rated impairments 416. Empirical experiments under varying conditions of signal power, noise level, and et cetera enable a transform 414 to be constructed which accurately predicts impairments 416 that will reported in response to a given combination of objective measurements 412. A manner in which transform 414 may be established by empirical experiments is described in U.S. Patent Application Serial No. 09/778,186, which is a continuation-in-part of U.S. Patent No. 6,370,120. Both U.S. Patent No. 6,370,120 and U.S. Patent Application Serial No. 09/778,186 are incorporated herein, by reference, in their entireties.

[0037] Transform 414 maps objective measurements 412 to impairments 416 regardless of the actual mechanism of signal transport chosen. Any communications service exhibiting a particular combination of signal and noise levels and other characteristics will be similarly rated by listeners regardless of whether the signal is actually traversing radio signals, PCM codecs, long-haul TDM equipment, or a packet-switched data network.

[0038] Where the quality of a particular communications service is to be measured, many sets of objective measurements 412 may be recorded. Each such set of

measurements is then converted using transform 414 into a set of impairments 416. A set of impairments 416 may comprise the four impairments LV, NS, EC, and SD each rated by test subjects as being “none”, “some” or “much”. It is also possible to directly collect subjective measurements of impairments 416 from a communications service by conducting test calls, although this practice is tedious and disadvantageous where transform 414 is already well established. As described above, measurement set 410 includes objective measurements 412 transformed into a set of impairments 416. In another embodiment, measurement set 410 includes a set of subjectively rated impairments 416 directly obtained from test calls through a communications service. A large number of measurement sets 410 pertaining to given communication service may be said to comprise the service attribute test (SAT) data 420 for the service. As described in U.S. Patent No. 6,370,120, the SAT data 420 for a given communication service may be convolved with an effects matrix 460 to yield an estimate of mean opinion score and P(UDI) for the service. Like transform 414, effects matrix 460 is a prediction model that may be constructed from large numbers of empirical observations to correlate reported impairments to overall MOS and P(UDI) values.

[0039] The relationships captured in the process displayed in Figure 4 enable analyses, conducted in the manner taught in U.S. Patent No. 6,370,120, to determine combinations of dropped frame rate (DFR) and round-trip packet-switched delay (RTL) that represent acceptable quality, relative to extant circuit switched telephony (e.g., wherein no packet-switching takes place). This is accomplished as follows:

[0040] A packet loss rate (PLR) transformation 422 is applied based upon a presumed packet loss rate 424. This transformation 422 redistributes, within SAT data 420, the proportions of reporting various combinations of impairments resulting, in this case, in the estimate of the speech distortion impairment, augmented in relation to how much packet loss is occurring. Transformation 422 draws upon a table 426 relating reports of speech distortion versus packet loss rate when all other impairments are negligible. The result of transformation 422 is transformed SAT data 430. SAT data 430 reflects the performance of a communication service suffering a level of packet loss rate in conjunction with whatever other impairments were present in SAT data 420 prior to PLR transformation 422.

[0041] Transformed SAT data 430 is then convolved with effects matrix 460 in process block 432. The result is an estimated set of MOS and P(UDI) values 440. These values take into account the rate of packet loss 424 provided as input earlier.

[0042] In step 444, the effects of packet delay are applied to the P(UDI). The effects of packet delay upon P(UDI) are calculated and applied to the P(UDI). It may be generally said that the greater the packet delay, the more P(UDI) is increased. It has been determined experimentally by the present inventors that packet delay affects P(UDI) independently of other impairments. Therefore, it suffices to apply the effects of packet delay after all other impairments are accounted for and to do so without repeating the development of transform 414 or effects matrix 460. After packet delay is taken into account in block 442, the result is the composite MOS and P(UDI) for the communication service with both a given packet delay and packet loss rate represented. The composite MOS and P[UDI] are then compared with limits corresponding to acceptable circuit switched telephony values, to produce a profile like that shown in Figure 5. Figure 5, of course, shows acceptable cumulative distribution functions (CDFs) for DFRs associated with the PLRs as a function of the RTL. This matrix is the reference matrix of the current invention.

[0043] As embodied herein and depicted in Figure 5, a matrix 500 illustrating the semantic content of the test matrix, the reference matrix and a reference pattern matrix is disclosed. Each of these are described in more detail below. As shown, the rows 502 are defined by variable round-trip latency intervals, in milliseconds. Columns 504 are defined in terms of variable dropped frame rate intervals, e.g., the average number of frames dropped per second. Thus, the boundaries defining row 502 and column 504 content can be readily changed. The intervals represent lower and upper values for a range of possible values that are used to produce counts of test data elements (in the test matrix) satisfying different conditions defined by round-trip latency (RTL) and dropped frame rates (DFR) for each recorded test.

[0044] A description of the test matrix is as follows. The test matrix is constructed by measuring a multiplicity of connections between a first location and a second location. The matrix is populated with a count of observed calls. For example, a five (5) in

element 506 indicates that five calls were observed having an RTL in the interval between 281-307msec, and a DFR less than 1.5.

[0045] The amount of data to be accumulated in populating the test matrices is user specified. This is implemented via one of three control options. In one embodiment, the number of data points to be accumulated in the test data matrix is specified before recording the test measurements. Alternatively, the number of scheduled probe tests to be accumulated in the test data matrix may be specified before recording. In the third alternative, the time period over which all available test results are to be accumulated is specified before the test data matrix is recorded. As they are created, the test results accumulated in any test data matrix must comprise only those tests run for one particular origin/destination pair (order considered). Data records comprising the content of a test result matrix shall be annotated to show: the probe from which the tests whose results are recorded in the test data matrix originated; the destination probe; the date and time of the start of the first test whose results are accumulated in the test data matrix; and the time of the end of the last test whose results are accumulated in the test data matrix.

[0046] As described above, the test data matrix is processed, in near real time, to determine the quality of voice over the packet-switched voice service. Interpretation of test data matrices to detect indications of unacceptable voice quality are accomplished in the following way. When the routine for processing the test data matrix, M_{ij} , is invoked, each row is summed step by step, to produce a cumulative matrix S_{ij} . S_{ij} is the sum of the M_{ik} values from $k = 1$ to $k = j$. For example, the third element in a particular row of cumulative matrix S_{ij} is the sum of the first, second, and third elements in the corresponding row of test data matrix M_{ij} . Likewise, the fourth element in a row of cumulative matrix S_{ij} is the sum of the first, second, third, and fourth elements in the corresponding row of test data matrix M_{ij} . This process is performed for each element in a row-by-row manner.

[0047] Each of the entries in each row of S_{ij} is divided by S_{in} , which is the last value in row S_i , to produce a cumulative distribution function (CDF) matrix P_{ij} . However, if the value of the last value (S_{in}) is zero, all the elements of the corresponding

row in P_{ij} are populated with zeroes. Otherwise, each row in P_{ij} will comprise a monotone non-decreasing sequence of values ending with the value 1.0.

[0048] As described above, the quality of service is determined by comparing the reference matrix to the test matrix. This step is performed indirectly. In actuality, CDF matrix P_{ij} is compared element-by-element to the like-sized reference matrix R_{ij} , which was described in detail above. This comparison produces a *reference pattern* matrix, RP_{ij} , whose entries are defined by the following criteria:

[0049] $RP_{ij} = 0$ for all values of j in a row i , for which the values P_{ij} are all zero. In other words, when the values of P_{ij} are zero, this indicates that no data was recorded, and the corresponding element of RP_{ij} is set to equal zero.

When $P_{ij} \neq 0$, then $RP_{ij} = 1$ if $P_{ij} < (R_{ij} - \Delta)$, and $RP_{ij} = 0$ if $P_{ij} \geq (R_{ij} - \Delta)$.

[0050] The variable " Δ " is used to reduce statistically insignificant false alarms produced by application of the routines that are described below. Delta " Δ " may be refined on the basis of experience. The value of Delta (Δ) is a user-specified variable or calculated function of the number of test calls represented in the test matrix. Its purpose is to set a threshold of significance of cell-by-cell differences as might, for example, be derived by application of the Kolmogorov/Smirnov for tests of significance of differences between CDFs.

[0051] The step of processing also includes the use of routines that are invoked to produce associated estimates of MOS and P[UDI] from a test data matrix. These routines are implemented by use of two auxiliary items. A MOS vector, VM; and a P[UDI] matrix, MP. For a given test data matrix, M, these items shall be used to calculate estimates of MOS and P[UDI] as follows. To determine MOS, the columns of M are added to produce the vector SC. Element-by-element multiplication of SC by VM is performed to produce the vector VS. To determine P[UDI], element-by-element multiplication of M by MP is performed to produce the matrix SP.

(1) $MOS = (\Sigma VS) / (\Sigma SC)$, wherein Σ denotes sum of all of the elements of

an array; and

$$(2) P[UDI] = (\Sigma SP)/(\Sigma M).$$

[0052] The above combination provides valuable information because, for example, where a MOS of 3.6 is very good and might normally be considered to guarantee user satisfaction, studies have shown that users will nonetheless complain of poor quality when the P(UDI) is greater than six percent regardless of the MOS. Therefore, predictions of overall service acceptability are best made considering both MOS and P(UDI).

[0053] Referring back to Figure 3, and in particular, to display 350, the present invention also includes software routines that are configured to automatically display a notice to maintenance personnel of the detection of a newly-generated reference pattern (RP) matrix containing at least one non-zero value. Such a notice shall take the form of a brief alarm message delivered to, or generated by, a host server disposed in the network maintenance facility, and written to a file of alarms. The reference pattern matrix contains enough data to enable data record parsing, to create a human readable amplifying display message. The alarm message is accompanied by a display of the reference pattern matrix and/or natural language descriptors of the conditions that can be inferred from the reference pattern matrix. The human readable message also includes a description of the origin/destination pair associated with the underlying test matrix, and the date and time of the creation of the test data matrix that generated the reference pattern matrix. Receipt of the brief alarm message shall generate a visible notice of this information on display 350.

[0054] Maintenance facility 30 and display 350 shown in Figure 3 include a graphical user interface (GUI). The GUI is equipped with routines that respond to a user selection of the alarm message described above. In one embodiment, the graphical user interface (GUI) is equipped to show a human readable description of the alarm detailing the associated origin/destination pair, and date and time of creation of the underlying test data matrix. A two dimensional display of the CDF matrix P_{ij} may also be displayed in a format like that shown in Figure 4. The matrix cells whose RP matrix values were non-

zero, are highlighted, displayed with a contrasting color, or otherwise visually distinguished. A line of information that displays the associated values of MOS and P[UDI] generated by invoking the routines described above may also be displayed by the GUI.

[0055] Referring to database 330 in Figure 3, the present invention also defines a compact data structure for collecting and archiving annotated test data matrices. Computer 340, or some other maintenance facility host, is configured to retrieve from database archives any test data matrices satisfying user-specified criteria for origin(s) of the test data, destination(s) of the test data, or date and/or time of collection. The data structure for each test matrix may include origination, destination, time/date of matrix creation, and a compact representation of the matrix itself. In one embodiment, the compact representation includes the position points of non-zero values in the matrix, and the values themselves.

[0056] Computer 340 is also configured to parse data retrieved from the archived annotated test data matrices stored in database 330. Computer 340 is also programmed to combine individual test data matrices to aggregate results. The aggregated results may be used by network analysts to detect service trends.

[0057] As embodied herein, and depicted in Figure 6, a block diagram of system 600 for evaluating the perceived quality of a packet switched voice channel in accordance with another embodiment of the present invention is disclosed. System 600 is configured to monitor the perceived quality of a call between customer transceiver 200 and customer transceiver 220 over packet transport network 18. Since transceiver 220 is identical to transceiver 200, only transceiver 200 will be discussed in detail. A network management system (NMS) 300 is coupled to packet switched network 18. Since we are mostly concerned with the receive side of transceiver 200, the discussion will center on that aspect of transceiver 200. Transceiver 200 (and 220) includes packet interface 202 which is coupled to packet transport network 18. Packet interface 202 processes the packet headers and provides codec 204 with a coded digital signal. Codec 204 decodes and de-compresses the coded signal and provides D/A converter 206 with a PCM digital signal. D/A converter 206 converts the digital signal into an analog electrical signal.

The analog electrical signal is converted into audible sound signal by a transducer, such as audio speaker 208.

[0058] In this embodiment, quality analyzer 610 may be coupled to packet interface 202 and/or the output of codec 204. Quality analyzer 610 is configured to analyze both the 64kbps digital signal coming from the codec, and/or packet transmission statistical data derived from packet interface 202. Each source may be employed to monitor the perceived quality of the packet voice transmission.

[0059] Those of ordinary skill in the art will recognize that quality analyzer 610 may be implemented in any number of ways. In one embodiment, quality analyzer 610 may be implemented in software resident in codec 204. The software may be executed by the processor in codec 204. In another embodiment, quality analyzer 610 is implemented as stand alone device similar to the one depicted in Figure 3.

[0060] In one embodiment, quality analyzer 610 is configured to analyze the digital output of the codec to obtain an instantaneous distortion measurement. The distortion measurement is reported to transceiver 220 via the far-end quality indicator 226 (Figure 6). In this embodiment, the distortion measurement need not be converted into a MOS or a P[UDI] to be reported. The distortion may be reported as a raw measurement, as a normalized score, such as a MOS or something similar thereto, or as a kurtosis value. Reference is made to U.S. Patent No. 6,246,978, which is incorporated herein by reference as though fully set forth in its entirety, for a more detailed explanation of obtaining distortion measurements directly from a digitized speech signal. In accordance with the teachings of the '978 patent, difference values may be calculated among successive samples in a digitized speech signal and the distribution of these signals may be analyzed. Abnormalities, such as kurtosis, in the distribution of first and/or second differences (which approximate first and second derivatives) may be correlated to perceived distortion.

[0061] Because distortion is such an important effect, in some instances it is an effect worth reporting by itself. As explained earlier, it is considered to be advantageous to obtain measurements from the output of the codec so that the actual impact of

transmission performance upon codec function may be taken into account. This approach is also advantageous because the arrangement of elements in Figure 6 and the function of quality analyzer 610 need not change as new codecs are developed and deployed in the future.

[0062] In another embodiment, quality analyzer 610 may be more sophisticated, having the same functionality provided by the system depicted in Figure 3. For example, quality analyzer 610 may be configured to obtain a set of measurements from the digital signal, and estimate the perceived level of impairment using the set of measurements. The set of impairments include low volume, noise, speech distortion, and echo. Analyzer 610 may be configured transform the impairment estimates into a Mean Opinion Score (MOS), and a probability that the calls would be considered as “unusable,” “difficult,” or “irritating” P[UDI].

[0063] As shown in Figure 6, analyzer 610 may also obtain packet transmission characteristics from packet interface 202. Packet transmission characteristics include packet delay, packet loss, and packet jitter. Packet transmission characteristics may be correlated with the MOS, P[UDI], and/or the perceived impairments to determine if a perceived quality problem is the result of a problem in packet transport network 18, or in a non-packet switched portion of network 10.

[0064] Measurements of quality obtained by analyzer 610 or quality indicators derived by analyzer 610 may be communicated to network management system 300 or to a far end transceiver 220 in the form of packetized data through network 18. Analyzer 610 may provide this data to packet interface 202 to be dispatched through network 18. Far-end quality indicator 226 may receive this data and provide a display, such as a bar graph, meter or numerical value to the far end user so that they know immediately how well they are being received. A far end user who observes a significant degradation while they are speaking may know to repeat certain portions of what they have said. Both participants may get a sense of the quality of the connection without having to ask one another if they are being received clearly.

[0065] It will be apparent to those skilled in the art that various modifications and variations can be made to the present invention without departing from the spirit and scope of the invention. Thus, it is intended that the present invention cover the modifications and variations of this invention provided they come within the scope of the appended claims and their equivalents.